

Efficient solutions for weight-balanced partitioning problems

Steffen Borgwardt

Technische Universität München

Shmuel Onn

Technion - Israel Institute of Technology, Haifa

Abstract

We prove polynomial-time solvability of a large class of clustering problems where a weighted set of items has to be partitioned into clusters with respect to some balancing constraints. The data points are weighted with respect to different features and the clusters adhere to given lower and upper bounds on the total weight of their points with respect to each of these features. Further the weight-contribution of a vector to a cluster can depend on the cluster it is assigned to. Our interest in these types of clustering problems is motivated by an application in land consolidation where the ability to perform this kind of balancing is crucial.

Our framework maximizes an objective function that is convex in the summed-up utility of the items in each cluster. Despite hardness of convex maximization and many related problems, for fixed dimension and number of clusters, we are able to show that our clustering model is solvable in time polynomial in the number of items if the weight-balancing restrictions are defined using vectors from a fixed, finite domain. We conclude our discussion with a new, efficient model and algorithm for land consolidation.

Keywords: constrained clustering, convex maximization, integer programming, land consolidation

2010 MSC: 90C10, 90C25, 90C90, 91C20

1. Introduction

Partitioning a set of items while respecting some constraints is a frequent task in exploratory data analysis, arising in both operations research and machine learning; see e.g. [3, 16]. We consider partitioning for which the sizes of the clusters are restricted with respect to multiple criteria. There are many applications where it is necessary to adhere to given bounds on the cluster sizes.

For example, these include the modeling of polycrystals in the material sciences [1] and face recognition using meshes, where the original mesh is partitioned into parts of equal sizes to obtain an optimal running time for graph-theoretical methods that are applied to all of these parts [4]. Our interest in these types of problems comes from an application

Email addresses: borgwardt@ma.tum.de (Steffen Borgwardt), onn@ie.technion.ac.il (Shmuel Onn)

in *land consolidation*. See for example [6, 8, 9] and in particular the outreach article [7] for the impact in academia and practice.

1.1. Land consolidation

The farmers of many agricultural and private forest regions in Bavaria and Middle Germany own a large number of small lots that are scattered over the whole region. The reasons are strict heritage laws and a frequent change of ownership. There is significant overhead driving and an unnecessarily high cost of cultivation. In such a situation, a land consolidation process may be initiated by the state to improve on the cost-effective structure of the region. Voluntary land exchanges (by means of lend-lease agreements) are a popular method for such a process: The existing lots are kept without changes and the rights of cultivation are redistributed among the farmers of the region.

This corresponds to a combinatorial redistribution of the lots and can be modeled as a clustering problem where each lot is an item and each farmer is a cluster [9]. The main goal is to create large connected pieces of land for each farmer. One way to do so is to represent the lots by their midpoints in the Euclidean plane and to use the geographical locations of the farmsteads of each farmer as a set of sites. Then one performs a weight-balanced least-squares assignment of the lots to these sites [5, 6, 11]. The result is a redistribution where the farmers' lots lie close to their farmsteads. As a positive sideeffect, many of a farmer's lots are connected and can be cultivated together.

The lots differ in several features like their size, quality of soil, shape, and attached subsidies, and some of these features are even different for each farmer. For example one farmer may be eligible for subsidies if they cultivate a given lot, while another farmer may not. In such a situation, the lot is more valuable for the first farmer. Here a natural constraint is that - after the redistribution - each farmer should have lots that (approximately) sum up to the farmer's original total with respect to each feature.

Of course, partitioning a weighted set of items into clusters of prescribed sizes (*weight-balanced partitioning*) is readily seen to be NP-hard, even for just two clusters and each item having just a single weight that is uniform for both clusters: deciding whether there is such a partition is at least as hard as the *Subset Sum* problem. The methods in the literature [5, 6, 7, 11] approach this intrinsic hardness by solving least-squares problems by an LP relaxation and rounding. The model in [6] performed particularly well in practice.

In this paper, we will present a general clustering framework that, when applied to land consolidation, improves on this model by dealing with its biggest shortcoming: The model in [6] (and in fact [5, 7, 11], too) is not able to balance the weights of clusters with respect to multiple features of the lots at the same time. Instead the redistribution of cultivation rights is done with respect to a single 'value' of a lot, an aggregation of all its properties (and this value is the same for each farmer).

The farmers will not accept a large deviation with respect to any of the features of their total lots. They will only participate in the redistribution if they do not lose a significant area of land, do not lose a lot in quality of soil, and do not lose much of their subsidies at the same time. But the models in the literature will return 'optimal' solutions for which not even the aggregated value necessarily is within the specified bounds (e.g. 3%) from the original. Further, even if a farmer's aggregated deviation is small, they may have received more land than before, but of much lower quality - which they will not accept. These are intrinsic weaknesses coming from the relaxation and rounding that are

performed in [5, 6, 7, 11]. In the practical implementation, this meant that a lot of the work still had to be done ‘by hand’.

However, the previous methods do not use all of the favorable properties of the input data for agricultural regions. The data typically falls into only a fixed number of categories, which just comes from the way the features are measured in practice. For example, especially large slots (> 5 hectare) are not traded at all. Further, one does not distinguish between lot sizes that differ by less than a tenth of a hectare, so that one obtains a finite domain of lot sizes. The quality of soil is measured with a number between 1 and 100, which is a finite domain itself. But in fact in a single agricultural region it is rare to have more than five different values within this range. The same happens for the subsidies attached to lots and other measures. With this additional assumption, we are able to present an algorithm that solves the problem of land consolidation exactly and efficiently.

1.2. Our contributions

In view of the above application, in the present paper we consider a new, generalized class of clustering problems where weight-balancing restrictions are defined using vectors from a fixed, finite domain. Each item specifies a vector of weights for each cluster. These weight vectors represent the weight the item would contribute with respect to the different features if assigned to the respective cluster. The sizes of the clusters are bounded above and below with respect to all features.

Further, each item gives a vector of utility values with respect to each cluster representing the utility gained if the item is assigned to the corresponding cluster. (These vectors do not have to be from a fixed domain.) Each cluster ‘collects’ a total utility by summing up the utility of all its items’ utility vectors. We then maximize an objective function that is convex on the sum of utility vectors of each part.

These are the functions used for shaped partitioning [2, 17, 18, 20], which encompass many of the objective functions commonly used in data analysis tasks, such as minimal-variance clustering; see e.g. [20]. They are intimately related to the studies of gravity polytopes [5, 11]. Shaped partitioning is known to be NP-hard if either the dimension of the utility vectors or the numbers of clusters is part of the input [17], because it captures the hardness of convex maximization due to the possible existence of exponentially many local optima [18]. For this, the number of clusters and the dimension of the utility vectors are fixed in our analysis.

Our main result is to show that this framework is polynomial-time solvable in the number of items, provided the weight vectors come from a fixed and finite domain. In Section 2, we introduce a formal notation for these clustering models, state our main results, and discuss the objective functions that can be represented in our models in some detail. Section 3 is dedicated to the necessary proofs. They are based on polynomial-size and -time reductions to the maximization of a (special) convex function over a system of constraints with an n -fold constraint matrix. The latter can be performed efficiently, which comes from a combination of recent algebraic Graver bases methods [12, 13, 18, 19] and geometric edge-directions and zonotope methods [17, 20]. We then transfer the complexity results to a generalized model for land consolidation, and present a new algorithm for it, in Section 4.

2. Model and results

In the following, we partition n items $\{1, \dots, n\}$ into p clusters π_1, \dots, π_p . Each item j has a utility matrix $C^j \in \mathbb{R}^{d \times p}$. Its i -th column C_i^j represents a vector of d utility values gained if item j is assigned to cluster π_i .

The utility of a clustering $\pi = (\pi_1, \dots, \pi_p)$ is $f(\sum_{j \in \pi_1} C_1^j, \dots, \sum_{j \in \pi_p} C_p^j)$, where $f : \mathbb{R}^{d \times p} \rightarrow \mathbb{R}$ is a convex function presented by a comparison oracle. Note that $\sum_{j \in \pi_i} C_i^j$ is the sum of utility vectors of the items in cluster π_i .

Recall that a comparison oracle for f is a routine that when queried on two values y, z returns whether $f(y) \leq f(z)$ or not. In our framework, we do not even need an explicit form of the function, only comparisons of its values. This makes the complexity results presented in the following stronger - essentially one works with as little information as possible. Of course, for any natural specific convex function that arises in data analysis, such as a norm or squared Euclidean distances (such as the ones that appear in Section 4), it is easy to explicitly perform these comparisons.

The task is to find a clustering of maximal utility under certain balancing constraints. We begin with a first model for these balancing constraints and extend it only later on. First, assume that we are given a fixed, finite set $\Omega \subset \mathbb{Z}^p$ such that each item j has a weight vector $w^j \in \Omega$. Informally, this vector contains a weight for each cluster π_i that it would contribute if assigned to π_i . Further, each cluster π_i defines a total weight $b_i \in \mathbb{Z}$ of the items that have to belong to it. Note that a finite cardinality m of Ω can e.g. be achieved by bounding the absolute values of all components in the w^j : for maximal absolute value ω , we obtain $m \leq (2\omega + 1)^p$.

The corresponding set of restrictions on the cluster sizes can be written as

$$\sum_{j \in \pi_i} w_i^j = b_i \quad (i \in [p]).$$

Let Π be the set of all partitions π of $\{1, \dots, n\}$ into p clusters. Then a full statement of this optimization problem (P_1) would be

$$(P_1) \quad \begin{aligned} \max_{\pi \in \Pi} \quad & f(\sum_{j \in \pi_1} C_1^j, \dots, \sum_{j \in \pi_p} C_p^j) \\ \sum_{j \in \pi_i} w_i^j = \quad & b_i \end{aligned} \quad (i \in [p])$$

Our first main result is polynomial-time solvability of this model for fixed d and p .

Theorem 1. *Suppose d and p are fixed and suppose there is a fixed, finite set $\Omega \subset \mathbb{Z}^p$ such that all $w^j \in \Omega$. Then for every convex f presented by a comparison oracle, problem (P_1) can be solved in polynomial time.*

Note that (P_1) and Theorem 1 already extend the model and complexity results in Section 3.3 in [13], where only the number of items in the clusters is balanced. So essentially the weights of all items are 1 for each cluster. In contrast, we allow for arbitrary weights of the items and these weights may also differ for the different clusters.

In many applications, instead of having exact sizes of the partition parts, one is given lower and upper bounds on the sizes. To extend the above program to lower and upper

bounds $b_i^\pm \in \mathbb{Z}$ on the total weights of the partition parts, we extend our formulation in several places. Formally, these constraints take the form

$$b_i^- \leq \sum_{j \in \pi_i} w_i^j \leq b_i^+ \quad (i \in [p]).$$

For our later proofs, we now rewrite the corresponding optimization problem using only equalities. Let us introduce slack variables $s_i^\pm \in \mathbb{Z}$ for all $i \in [p]$. Then the new optimization problem (P_2) can be stated as

$$(P_2) \quad \begin{aligned} \max_{\pi \in \Pi} \quad & f\left(\sum_{j \in \pi_1} C_1^j, \dots, \sum_{j \in \pi_p} C_p^j\right) \\ \left(\sum_{j \in \pi_i} w_i^j\right) + s_i^+ &= b_i^+ & (i \in [p]) \\ \left(\sum_{j \in \pi_i} w_i^j\right) - s_i^- &= b_i^- & (i \in [p]) \\ s_i^\pm &\geq 0 & (i \in [p]) \end{aligned}$$

With some modifications to the construction for (P_1) in the proof of Theorem 1, we can show a similar statement for this more general class of problems.

Theorem 2. *Suppose d and p are fixed and suppose there is a fixed, finite set $\Omega \subset \mathbb{Z}^p$ such that all $w^j \in \Omega$. Then for every convex f presented by a comparison oracle, problem (P_2) can be solved in polynomial time.*

Finally, we extend our model to allow for balanced weights with respect to $s \geq 1$ different features. Instead of item j listing a vector $w^j \in \mathbb{Z}^p$ of weights with respect to the clusters, it now has a matrix $W^j \in \mathbb{Z}^{s \times p}$ listing s -dimensional vectors W_i^j of weights contributed to the cluster π_i it is assigned to. All of these weight matrices come from a fixed, finite set $\Omega \subset \mathbb{Z}^{s \times p}$. As before, a finite cardinality m of Ω can be achieved by bounding the absolute values of all components in the W^j : for maximal absolute value ω , one now obtains $m \leq (2\omega + 1)^{sp}$.

Further, instead of $b_i^\pm \in \mathbb{Z}$, we now use $B_i^\pm \in \mathbb{Z}^s$, and likewise we have to use slack vectors $S_i^\pm \in \mathbb{Z}^s$ for all $i \in [p]$. This gives us an optimization problem (P_3) as

$$(P_3) \quad \begin{aligned} \max_{\pi \in \Pi} \quad & f\left(\sum_{j \in \pi_1} C_1^j, \dots, \sum_{j \in \pi_p} C_p^j\right) \\ \left(\sum_{j \in \pi_i} W_i^j\right) + S_i^+ &= B_i^+ & (i \in [p]) \\ \left(\sum_{j \in \pi_i} W_i^j\right) - S_i^- &= B_i^- & (i \in [p]) \\ S_i^\pm &\geq 0 & (i \in [p]) \end{aligned}$$

Even for this extension, we keep polynomial-time solvability.

Theorem 3. *Suppose d and p are fixed and suppose there is a fixed, finite set $\Omega \subset \mathbb{Z}^{s \times p}$ such that all $W^j \in \Omega$. Then for every convex f presented by a comparison oracle, problem (P_3) can be solved in polynomial time.*

We conclude this section with some general remarks on the expressive power of the objective function. In general, the convex functions $f : \mathbb{R}^{d \times p} \rightarrow \mathbb{R}$ of the form $f\left(\sum_{j \in \pi_1} C_1^j, \dots, \sum_{j \in \pi_p} C_p^j\right)$

cover a wide range of objective functions commonly used in clustering. In particular, they directly represent the *aggregation of utility values*. Each of the p clusters π_i contributes a summed-up utility vector $\sum_{j \in \pi_i} C_i^j$ of its items and these vectors then are aggregated to a value in \mathbb{R} by the convex function f . A simplest case for f is to just add up the components of all these vectors possibly (scaled by a factor). If the components of all C_i^j are non-negative, such an objective function corresponds to a linear transform of the l_1 -norm. Recall that all norms are convex functions. By choosing l_s with $s > 1$, one values single components in the utility vectors of the clusters higher in relation to many equally large values. For example, for l_∞ , only the largest absolute values among all components counts.

By means of our framework, one can represent even more general ways of aggregating the utility by turning to the so-called *clustering bodies* [10]: Here one combines two norms, one $\|\cdot\|$ for \mathbb{R}^d on the parts $\sum_{j \in \pi_i} C_i^j$ and a monotone norm $\|\cdot\|_*$ for \mathbb{R}^p aggregating these values. The objective function then takes the form

$$f\left(\sum_{j \in \pi_1} C_1^j, \dots, \sum_{j \in \pi_p} C_p^j\right) = \left\| \left(\left\| \sum_{j \in \pi_1} C_1^j \right\|, \dots, \left\| \sum_{j \in \pi_p} C_p^j \right\| \right) \right\|_*.$$

The level set for value at most one is a convex body, a clustering body, and serves as the unit ball for a semi-norm. This explains convexity of the above f .

3. Proofs

The proofs for Theorems 1 to 3 have a common demeanor in that we exhibit polynomial-time and -size transformation of the corresponding model to a variant of shaped partitioning for a set of constraints defined by an n -fold matrix. Generally speaking, we aim for a problem statement of the form

$$\max\{f(Cx) : A^{(n)}x = b, l \leq x \leq u, x \in \mathbb{Z}^N\} \quad (1)$$

with $l, u \in \mathbb{Z}^N$, $C \in \mathbb{Z}^{c \times N}$, and a convex function $f : \mathbb{Z}^c \rightarrow \mathbb{R}$ presented by a comparison oracle. Further, the n -fold matrix $A^{(n)}$ is derived from two matrices A_1 and A_2 by the classical construction in the form

$$A^{(n)} = \begin{bmatrix} A_1 & \dots & A_1 \\ A_2 & & \\ & \dots & \\ & & A_2 \end{bmatrix}$$

with n copies of both A_1 and A_2 arranged in the depicted layout. Note that n corresponds to the number of items; each of them gets a set of columns with one of the building blocks A_1 and A_2 . Clearly, N has to be a multiple of n .

A problem of the form (1) is polynomial-time solvable if A_1 and A_2 are fixed and C has a constant number of rows c . The following proposition sums up this important tool for our proofs.

Proposition 1 ([13]). *Let $c \in \mathbb{N}$ be fixed, and let A_1, A_2 be fixed matrices. Further, let $A^{(n)}$ be the n -fold matrix derived from A_1 and A_2 , let $l, u \in \mathbb{Z}^N$, let $C \in \mathbb{Z}^{c \times N}$, and $f : \mathbb{Z}^c \rightarrow \mathbb{R}$ presented by a comparison oracle.*

Then there is an algorithm that is polynomial in n and the length of binary input for C, l, u and b that solves

$$\max\{f(Cx) : A^{(n)}x = b, l \leq x \leq u, x \in \mathbb{Z}^N\}.$$

In view of Proposition 1, to obtain polynomial-time-solvability for our models, we will show that there is a polynomial-time and -size reduction to the above form in the size of the input for our original models. Note that a representation of the data as above is in stark contrast to the ‘natural’ representation in Section 2. In addition to the explicit prerequisites of Proposition 1, namely the constant number of rows for C and the fixed A_1, A_2 , we also have to guarantee that N stays polynomial in the size of the original input.

In the following, we carefully perform the necessary reductions. We begin with the first model (P_1) and Theorem 1, and then subsequently extend the constructions.

Proof of Theorem 1. We prove the claim by transforming (P_1) to a statement in the form

$$\max\{f(Cx) : A^{(n)}x = b', l \leq x \leq u, x \in \mathbb{Z}^N\}.$$

Let us begin with the constraints. Recall that m is the size of Ω . Consider the two matrices

$$A_1 = \begin{bmatrix} w_1^1 & & & & w_1^m & & \\ & w_2^1 & & \dots & & w_2^m & \\ & & \dots & & & & \dots \\ & & & w_p^1 & & & \\ & & & & & & w_p^m \end{bmatrix} \in \mathbb{R}^{p \times mp}$$

and

$$A_2 = [1 \ 1 \ \dots \ 1] \in \mathbb{R}^{1 \times mp}.$$

Further, we define the vectors $b = (b_1, \dots, b_p)^T$ and $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^n$ to be able to write a system of equations $A^{(n)}x = b' = \begin{pmatrix} b \\ \mathbf{1} \end{pmatrix}$. The variables $x \in \mathbb{R}^{n(mp)}$ here correspond to decision variables in the following way:

We have $x = (x^1, \dots, x^n)^T$, where x^j corresponds to the j -th column-block of the n -fold matrix, i.e. to the j -th ‘copy’ of A_1 and A_2 . Each x^j takes the form $x^j = (x_1^j, \dots, x_m^j)^T$, i.e. it consists of m blocks $x_i^j \in \mathbb{Z}^p$. Note that this construction yields $N = n(mp)$, which is polynomial in the input.

Further, let us define lower and upper bounds on x in the form $l \leq x \leq u$: Choose lower bounds $l = \mathbf{0} \in \mathbb{R}^N$ and define the upper bound vector $u = (u^1, \dots, u^n)^T$ to consist of n blocks $u^j = (u_1^j, \dots, u_m^j)^T \in \mathbb{Z}^{mp}$, setting $u_i^j = \mathbf{1} \in \mathbb{Z}^p$ if item j has weight vector w^i and $u_i^j = \mathbf{0} \in \mathbb{Z}^p$ otherwise.

The system

$$A^{(n)}x = \begin{pmatrix} b \\ \mathbf{1} \end{pmatrix}, \quad l \leq x \leq u$$

can be derived in polynomial time and is of polynomial size, as it only uses a polynomial number of copies of numbers from the original input (and of zeroes and ones). Let us discuss why it is equivalent to our original set of constraints. First, note that the lower and upper bounds force an integral solution x to be a 0, 1-solution.

The assignment of item j to a cluster is determined by the decision variables x^j corresponding to the j -th block of A_1 and A_2 . The ‘ A_2 -block’ of the system tells us that

precisely one entry in x^j is equal to 1. By the upper bounds on x , this can only be the case for an index which corresponds to a correct combination of cluster π_i and weight contribution by item j . For this, in the ' A_1 -block' of the system, the correct weight for x_j is added up in the equation to obtain total cluster size b_i . Thus all items are assigned and all clusters obtain the correct total weight.

It remains to check whether the objective function can be written in form $f(Cx)$, where C has a constant number of rows. Note that

$$\sum_{j \in \pi_s} C_s^j = \sum_{j=1}^n \sum_{i=1}^m ((x_i^j)^T \mathbf{1}) C_s^j.$$

Thus C takes the form

$$C = [(C')^1 \dots (C')^1 \dots (C')^n \dots (C')^n] \in \mathbb{R}^{dp \times n(mp)},$$

with m consecutive copies for each of the $(C')^j$. These are defined as

$$(C')^j = \begin{bmatrix} C_1^j & & & \\ & C_2^j & & \\ & & \dots & \\ & & & C_p^j \end{bmatrix} \in \mathbb{R}^{dp \times p}.$$

The claim now follows from the number of dp rows being constant and observing that Cx yields the vector $((\sum_{j \in \pi_1} C_1^j)^T, \dots, (\sum_{j \in \pi_p} C_p^j)^T)^T$. We satisfy all prerequisites of Proposition 1 and are done. \square

Next, we prove Theorem 2 by extending the above construction. The goal is to enter the slack variables s_i^\pm into the model while preserving the n -fold structure of the program. It is possible to do so by duplicating them and moving them into the building blocks of the matrix as follows.

Proof of Theorem 2. Let A, b etc. refer to the construction for Theorem 1. Their corresponding new counterparts are denoted by a bar, e.g. \bar{A}, \bar{b} . We now want a problem statement for (P_2) of the form

$$\max\{f(\bar{C}x) : \bar{A}^{(n)}\bar{x} = \bar{b}', \bar{l} \leq \bar{x} \leq \bar{u}, \bar{x} \in \mathbb{Z}^{\bar{N}}\},$$

where $\bar{A}^{(n)}$ is an n -fold matrix derived by the standard construction using matrices \bar{A}_1 and \bar{A}_2 .

For this formulation, we use a vector $\bar{x} = (\bar{x}^1, \dots, \bar{x}^n)^T$, where the column-blocks $\bar{x}^j = ((s^j)^+, (s^j)^-, x^j)^T$ now also have copies of the slack variables $(s^j)^+ = ((s_1^j)^+, \dots, (s_p^j)^+)^T$ and $(s^j)^- = ((s_1^j)^-, \dots, (s_p^j)^-)^T$ for each of the original x^j . Note that $\bar{x} \in \mathbb{R}^{\bar{N}}$, where $\bar{N} = n(2p + mp)$, which again is polynomial in the input.

$\bar{A}^{(n)}$ is set to

$$\bar{A}^{(n)} = \begin{bmatrix} \bar{A}_1 & \dots & \bar{A}_1 \\ \bar{A}_2 & & \\ & \dots & \\ & & \bar{A}_2 \end{bmatrix},$$

where

$$\bar{A}_1 = \begin{bmatrix} E & \mathbf{0} & A_1 \\ \mathbf{0} & -E & A_1 \end{bmatrix} \in \mathbb{R}^{2p \times (2p+mp)} \text{ and } \bar{A}_2 = \begin{bmatrix} 0 & 0 & A_2 \end{bmatrix} \in \mathbb{R}^{1 \times (2p+mp)}$$

with $E \in \mathbb{R}^{p \times p}$ the unit-matrix and $\mathbf{0} \in \mathbb{R}^{p \times p}$ the 0-matrix. Further $\bar{b}' = (\bar{b}, \mathbf{1})^T$ with $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^n$ and $\bar{b} = (b_1^+, \dots, b_p^+, b_1^-, \dots, b_p^-)^T$ listing the upper and lower bounds on the cluster sizes.¹

Next, let us define lower and upper bounds $\bar{l}, \bar{u} \in \mathbb{R}^{\bar{N}}$ for \bar{x} . As before, $\bar{l} = \mathbf{0} \in \mathbb{R}^{\bar{N}}$. Further, we use the upper bound vector $\bar{u} = (\bar{u}^1, \dots, \bar{u}^n)^T$, where $\bar{u}^1 = (\nu, \dots, \nu, u^1)^T \in \mathbb{R}^{2p+mp}$ begins with $2p$ entries of value $\nu = \sum_{i=1}^p \sum_{j=1}^n |w_i^j|$. Further $\bar{u}^i = (0, \dots, 0, u^i)^T \in \mathbb{R}^{2p+mp}$ begins with $2p$ entries 0 for all $i > 1$. Note that in a feasible solution no slack variable ever is larger than ν and that ν has an encoding length that is polynomial in the input size of (P_2) .

Finally, we construct the new matrix $\bar{C} = (\bar{C}^1, \dots, \bar{C}^n) \in \mathbb{R}^{dp \times N}$ of n building-blocks $\bar{C}^j = [\mathbf{0} \ (C')^j \dots (C')^j] \in \mathbb{R}^{dp \times 2p+mp}$, where $\mathbf{0} \in \mathbb{R}^{dp \times 2p}$ is a 0-matrix and there again are m copies of the $(C')^j$ in the blocks \bar{C}^j . Essentially, all parts that correspond to slack variables are ignored for the objective function value, i.e. $f(\bar{C}\bar{x}) = f(Cx)$.

The new system $\bar{A}^{(n)}\bar{x} = \bar{b}', \bar{l} \leq \bar{x} \leq \bar{u}$, as well as the new objective function matrix \bar{C} , extend the original system only by introducing an additional polynomial number of zeroes, ones and copies of numbers from the input. The matrices \bar{A}_1, \bar{A}_2 still have a constant number of rows and columns, and \bar{C} has a constant number of rows.

It remains to discuss why the new system is equivalent to the constraints of (P_2) . Clearly, the ' x -part' of \bar{x} still has to be a 0,1-solution. These 0,1-entries have the same role as in the original construction, as in the \bar{A}_2 -part the slack variables are only combined with zeroes.

The lower bounds on the slack variables force them to be non-negative. By the upper bounds, only the slack variables in the first block of the construction can be greater than zero. The other slack variables are only in the system to preserve the n -fold structure of the matrix. Note that it suffices to allow integral values for the slack variables (recall $\bar{x} \in \mathbb{Z}^{\bar{N}}$), as both the w_i^j and the b_i^\pm are integral. The rows of the \bar{A}_1 -blocks then guarantee that the cluster sizes plus the slack variables add up to the given lower and upper bounds. In particular, the clusters' sizes lie in the given, bounded range.

Thus the new system corresponds to (P_2) and satisfies all prerequisites of Proposition 1, in particular \bar{A}_1, \bar{A}_2 , and \bar{C} do so. \square

Finally, we prove Theorem 3 by extending the construction in the proof of Theorem 2 to allow for s different features for each item. We continue with the notation introduced for the proofs of both Theorem 1 and 2.

Proof of Theorem 3. We want a problem statement of (P_3) of the form

$$\max\{f(\tilde{C}\tilde{x}) : \tilde{A}^{(n)}\tilde{x} = \tilde{b}', \tilde{l} \leq \tilde{x} \leq \tilde{u}, \tilde{x} \in \mathbb{Z}^{\tilde{N}}\}, \quad (P^*)$$

¹In the notation of \bar{b}' and in other places, we avoid double transposes when writing vectors for better readability when the context is clear. For example $c = (a, b)^T$ for two column vectors a, b would be a column vector c .

where $\tilde{A}^{(n)}$ is an n -fold matrix derived by the standard construction for matrices \tilde{A}_1 and \tilde{A}_2 .

We use $\tilde{x} = (\tilde{x}^1, \dots, \tilde{x}^n)^T$, where the column-blocks $\tilde{x}^j = ((S^j)^+, (S^j)^-, x^j)^T$ all contain their own copies of the slack vectors $(S^j)^+ = ((S_1^j)^+, \dots, (S_p^j)^+)^T$ and $(S^j)^- = ((S_1^j)^-, \dots, (S_p^j)^-)^T$. Herewith $\tilde{x} \in \mathbb{R}^{\tilde{N}}$ for $\tilde{N} = n(2sp + mp)$, which is polynomial in the size of the input.

For the n -fold construction of $\tilde{A}^{(n)}$, let now

$$\tilde{A}_1 = \begin{bmatrix} E & \mathbf{0} & A'_1 \\ \mathbf{0} & -E & A'_1 \end{bmatrix} \in \mathbb{R}^{2sp \times (2sp+mp)} \text{ with}$$

$$A'_1 = \begin{bmatrix} W_1^1 & & & & W_1^m & & \\ & W_2^1 & & \dots & & W_2^m & \\ & & \dots & & & & \dots \\ & & & W_p^1 & & & \\ & & & & & & W_p^m \end{bmatrix} \in \mathbb{R}^{sp \times mp},$$

where $E \in \mathbb{R}^{sp \times sp}$ is the unit-matrix and $\mathbf{0} \in \mathbb{R}^{sp \times sp}$ the 0-matrix. Further, let

$$\tilde{A}_2 = \begin{bmatrix} 0 & 0 & A_2 \end{bmatrix} \in \mathbb{R}^{1 \times (2sp+mp)}.$$

The new right-hand side vector is $\tilde{b}' = (\tilde{b}, \mathbf{1})^T \in \mathbb{R}^{2sp+n}$ with $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^n$ and $\tilde{b} = (B_1^+, \dots, B_p^+, B_1^-, \dots, B_p^-)^T$ listing the upper and lower bound vectors on the cluster sizes.

Again, we use $\tilde{l} = \mathbf{0} \in \mathbb{R}^{\tilde{N}}$ as lower bounds. The upper bound vector $\tilde{u} = (\tilde{u}^1, \dots, \tilde{u}^n)^T \in \mathbb{R}^{\tilde{N}}$ contains $\tilde{u}^1 = (\tilde{v}, \dots, \tilde{v}, u^1)^T \in \mathbb{R}^{2sp+mp}$ which begins with $2p$ vectors $\tilde{v} = \sum_{i=1}^p \sum_{j=1}^n |W|_i^j \in$

\mathbb{R}^s (where $|W|_i^j$ refers to the vector listing the absolutes in W_i^j componentwisely), and $\tilde{u}^i = (0, \dots, 0, u^i)^T \in \mathbb{R}^{2sp+mp}$ begins with $2sp$ entries 0 for all $i > 1$. The $u^j = (u_1^j, \dots, u_m^j)^T$ are defined by setting $u_i^j = 1 \in \mathbb{Z}^p$ if item j has weight matrix W^i and $u_i^j = 0 \in \mathbb{Z}^p$ otherwise.

Finally, as in the construction for Theorem 2, it is necessary to ignore the parts that correspond to slack variables for the objective function value. It is possible to do so by means of the matrix $\tilde{C} = (\tilde{C}^1, \dots, \tilde{C}^n) \in \mathbb{R}^{dp \times \tilde{N}}$ that consists of n building-blocks $\tilde{C}^i = [\mathbf{0} \ (C')^1 \dots (C')^1] \in \mathbb{R}^{dp \times 2sp+mp}$, where $\mathbf{0} \in \mathbb{R}^{dp \times 2sp}$ is a 0-matrix. Then $f(\tilde{C}\tilde{x}) = f(\tilde{C}\tilde{x}) = f(Cx)$.

It remains to explain why this formulation represents the constraints of (P_3) . By definition of \tilde{l} and \tilde{u} , the 'x-part' of \tilde{x} still is a 0, 1-solution. These 0, 1-entries play the same role as in the original construction by definition of the \tilde{A}_2 . The slack variables are non-negative, and by the upper bounds only the slack variables in the first block of the construction can be greater than zero. The components of \tilde{v} are sufficiently large to not impose a restriction. Again it suffices to allow integral values for the slack variables, as both the W^j and the B_i^\pm are integral. The rows of the \tilde{A}_1 -blocks then guarantee that the cluster sizes plus the slack variables add up to the given lower and upper bounds for each component of the weight vectors of the clusters, which implies that all of the clusters' total weights lie within the range given by the B_i^\pm .

Thus the new system corresponds to (P_3) and satisfies all prerequisites of Proposition 1. \square

4. Polynomial-time land consolidation

Our interest in the presented clustering framework whose favorable complexity we studied in Sections 2 and 3 originally came from an application in land consolidation. We now conclude our discussion with a new model and algorithm for it. Recall the description of the problem in Section 1.1.

First, let us connect the general definitions used for model (P_3) with the application. We here extend [6], which was a particularly useful model in practice. The items are the n lots which have to be divided among the p farmers π_1, \dots, π_p (the clusters). We begin with the balancing constraints.

4.1. Weight-balancing constraints

The lots differ in s different features, for example size, value (which is impacted by the quality of soil), attached subsidies, and so on. While some of these measures are independent of which farmer the lot is assigned to - for example the size and quality of soil are just fixed numbers - for others this may not be the case. The subsidies a farmer gets for cultivating a lot depend on several factors that differ between farmers - for example depending whether the farmer represents a small local family or a large agricultural business working in multiple regions.

For each lot, we set up a matrix $W^j \in \mathbb{R}^{s \times p}$. The column vectors W_i^j list the contribution of lot j to the total of a farmer π_i with respect to the s features, if the lot is assigned to the farmer. The objective functions in the following will use the sizes of the lots of each farmer. The information on the size of lots is represented as one of the components in each column W_i^j of $W^j \in \mathbb{R}^{s \times p}$. We give the size of lot j the explicit name ω_j .

Let $B_i \in \mathbb{R}^s$ list the summed-up total features for farmer π_i . In the redistribution process, the farmers do not accept a large deviation with respect to any of the features. For an accepted change of for example 3%, one obtains $B_i^- = 0.97 \cdot B_i$ and $B_i^+ = 1.03 \cdot B_i$. The accepted deviations may also be defined differently for the different features.

4.2. Objective function

In the following, in a generalization of the model in [6], we discuss a family of objective functions for land consolidation that fit with our framework. Each lot j is represented by its midpoint $z_j \in \mathbb{R}^2$. Further, the farmers π_i specify a location $v_i \in \mathbb{R}^2$ of their farmstead. The distance of a lot and a farmstead is measured by the square of their Euclidean distance $\|v_i - z_j\|^2$.

We begin with the classical least-squares assignment for a given, single location v_i of a farmstead (as a basic building block of our model). The lots are assigned to the farmers such that

$$\sum_{i=1}^p \sum_{j \in \pi_i} \omega_j \|v_i - z_j\|^2 \quad (f_1)$$

is minimized. The sizes ω_j of the lots are used as scaling factors for the distances $\|v_i - z_j\|^2$ to have a fair treatment of the assignment of one large lot or of many small lots.

It remains to explain why this objective function fits our framework, i.e. that it can be represented by a function $f : \mathbb{R}^{d \times p} \rightarrow \mathbb{R}$ that is passed the p arguments of the form $\sum_{j \in \pi_i} C_i^j$ and is convex on each of these sums [20]:

For each item j , one can use $C_i^j = -\omega_j \|v_i - z_j\|^2$, so that $\sum_{j \in \pi_i} C_i^j = -\sum_{j \in \pi_i} \omega_j \|v_i - z_j\|^2$.

It then suffices to maximize the linear function f that sums up its arguments. Formally $f(y) = \mathbf{1}^T y$ with $\mathbf{1} = (1, \dots, 1)^T \in \mathbb{R}^p$. Note $d = 1$.

Such a least-squares assignment favors a good assignment of lots of a farmer with larger total size over a good assignment of a smaller farmer. Thus one may want to 'normalize' the different parts of this sum. We introduce this new approach as **normed least-squares assignment**.

For this, define $C_i^j = \left(\frac{-\omega_j \|v_i - z_j\|^2}{\omega_j} \right) \in \mathbb{R}^2$ and use a function $f : \mathbb{R}^{2 \times p} \rightarrow \mathbb{R}$ that sends $y = \left((y_{11}, y_{12}), \dots, (y_{p1}, y_{p2}) \right)$ to $f(y) = \sum_{i=1}^p \frac{y_{i1}}{y_{i2}}$. As $\sum_{j \in \pi_i} C_i^j = \left(-\frac{\sum_{j \in \pi_i} \omega_j \|v_i - z_j\|^2}{\sum_{j \in \pi_i} \omega_j} \right)$, one obtains a maximization of

$$-\sum_{i=1}^p \frac{1}{\sum_{j \in \pi_i} \omega_j} \sum_{j \in \pi_i} \omega_j \|v_i - z_j\|^2,$$

or equivalently, and more intuitively, a minimization of

$$\sum_{i=1}^p \frac{1}{\sum_{j \in \pi_i} \omega_j} \sum_{j \in \pi_i} \omega_j \|v_i - z_j\|^2. \quad (f_2)$$

By scaling by the inverse of each farmer's total assigned land, each farmer's quality of assignment contributes equally to the final objective function value.

Unfortunately, $f(y) = \sum_{i=1}^p \frac{y_{i1}}{y_{i2}}$ is not convex, even when restricted to a strictly positive domain, so it is necessary to resort to a (provably good) approximation for f_2 . For this, denote the original total size of each farmers' lots by κ_i and the accepted lower and upper bounds on the sizes as κ^\pm . One can then use these apriori κ_i to estimate the total size of lots of each farmer after the redistribution. (Compare this to the use of an approximate center of gravity in [11].) This results in a minimization of

$$\sum_{i=1}^p \frac{1}{\kappa_i} \sum_{j \in \pi_i} \omega_j \|v_i - z_j\|^2, \quad (f_3)$$

which is a linear transform of a least-squares assignment and thus fits our framework, too. One can easily see this by using $C_i^j = -\frac{\omega_j}{\kappa_i} \|v_i - z_j\|^2$. The approximation error of optimizing f_3 in place of f_2 is provably low.

Lemma 1. *Let π be an optimal partition for f_3 and π' be optimal for f_2 . Then π is a $(\max_{i \leq p} \frac{\kappa_i^-}{\kappa_i}) (\max_{i \leq p} \frac{\kappa_i^+}{\kappa_i})$ -approximation with respect to f_2 .*

Proof. For a simple notation, we refer to the corresponding objective function values as $f_3(\pi)$ and $f_2(\pi')$. As $\kappa_i^- \leq \sum_{j \in \pi_i} \omega_j \leq \kappa_i^+$, one has both $f_2(\pi) \leq (\max_{i \leq p} \frac{\kappa_i^-}{\kappa_i}) f_3(\pi)$ and

$$f_3(\pi) \leq f_3(\pi') \leq (\max_{i \leq p} \frac{\kappa_i^+}{\kappa_i}) f_2(\pi'), \text{ which combines to } f_2(\pi) \leq (\max_{i \leq p} \frac{\kappa_i^-}{\kappa_i}) \cdot (\max_{i \leq p} \frac{\kappa_i^+}{\kappa_i}) \cdot f_2(\pi').$$

Thus we obtain a $(\max_{i \leq p} \frac{\kappa_i^-}{\kappa_i}) (\max_{i \leq p} \frac{\kappa_i^+}{\kappa_i})$ -approximation. \square

For an accepted upper and lower deviation of 3%, f_3 would yield a 1.0609-approximation error for f_2 . Note that f_3 is linear and is defined with $d = 1$, but uses the fact that $C_{i_1}^j$ and $C_{i_2}^j$ can differ for $i_1 \neq i_2$. In the literature there are several examples for objective functions that use a larger d , but that only fit our framework for $\kappa_i^- = \kappa_i^+$ for all $i \leq p$. Examples include finding a partition of minimal variance [13, 20] and pushing apart the centers of gravity of the partition parts [7, 10]. Both of them can then be interpreted as norm-maximization over a gravity polytope or a shaped partition polytope; recall the final remarks in Section 2.

4.3. Algorithm and Efficiency

Theorem 3 implies that for a fixed number p of farmers, and if there is a fixed set Ω of vectors of lot features, the above model is solvable in polynomial time for (f_1) and (f_3) in the size of the input, in particular n . (Recall that due to the way the features of the lots are measured, this is not a particularly restrictive assumption.) In combination with Lemma 1, we also obtain a polynomial-time approximation for (f_2) . Let us sum up this information.

Theorem 4. *Suppose the number of farmers p is fixed and suppose there is a fixed set $\Omega \subset \mathbb{Z}^{s \times p}$ of weight matrices such that all $W^j \in \Omega$. Then the above model for land consolidation can be solved exactly in polynomial time for objective functions f_1 and f_3 . For f_2 a $(\max_{i \leq p} \frac{\kappa_i}{\kappa_i^-})(\max_{i \leq p} \frac{\kappa_i^+}{\kappa_i})$ -approximation can be computed in polynomial time.*

Algorithm 1 describes the method in pseudocode for objective function (f_3) ; we use the notation from the proof of Theorem 3. First, the input data is used for a formulation (P^*) of the problem in the form

$$\max\{f(\tilde{C}\tilde{x}) : \tilde{A}^{(n)}\tilde{x} = \tilde{b}', 0 \leq \tilde{x} \leq \tilde{u}, \tilde{x} \in \mathbb{Z}^{\tilde{N}}\}. \quad (P^*)$$

For better readability, we have split up the description in two parts: First, one sets up a problem statement in the form (P_3) (step 1), then this formulation is transformed to the form (P^*) (step 2). In our practical implementation, we of course construct the matrices and vectors for (P^*) directly. In Section 3, we showed that this can be done in polynomial time and that (P^*) is of polynomial size in the input.

The input to the problem is an actual distribution of farmland, which we represent by $y \in \{0, 1\}^{p \times n}$ with $y_{ij} = 1$ if lot j belongs to farmer i and $y_{ij} = 0$ otherwise. In particular, the original distribution is feasible with respect to the given lower and upper bounds B_i^\pm . Thus there always is a feasible solution to (P_3) , and thus to (P^*) , as well. The proofs of Theorem 1 and 3 explain how to construct an initial feasible solution \tilde{x}_* for (P^*) from y . This is done in step 3; the original vector is ‘expanded’ by a polynomial number of zeroes in the correct components. Analogously, in the final step 5 of the algorithm, it then is easy to return an optimal vector of decision variables $y^* \in \{0, 1\}^{p \times n}$ from the optimal solution \tilde{x}^* for (P^*) .

Steps 1 to 3 provide a suitable formulation (P^*) and an initial feasible solution, such that the problem can be solved efficiently with tools from the literature. In step 4, one iteratively augments this initial feasible solution to an optimal one using Graver bases methods. These methods were first shown to run in polynomial time in [14] and drastically improved in [15] which is the fastest algorithm to-date.

As a service to the reader, let us briefly explain the main ingredients of this approach. For a more detailed background see the survey [19] and the books [12, 18], as well as the papers [13, 14, 15]. The basic notion underlying these methods is the so-called *Graver basis* $\mathcal{G}(\tilde{A}^{(n)})$ of the matrix $\tilde{A}^{(n)}$. Informally, this is a finite set of vectors g with some special properties: in particular, if a current solution \tilde{x} is not optimal, and f is linear, there is a $g \in \mathcal{G}(\tilde{A}^{(n)})$ and an $\alpha \in \mathbb{N}$ such that $\tilde{x} + \alpha \cdot g$ is another feasible solution and satisfies $f(\tilde{x}) < f(\tilde{x} + \alpha \cdot g)$. This gives rise to an iterative augmentation scheme, which is performed in step 4.

Recall that Ω is fixed in our setting. In turn, this implies that the matrices \tilde{A}_1 and \tilde{A}_2 are fixed (the proof of this fact was one of our main concerns in Section 3). Thus the Graver basis $\mathcal{G}(\begin{pmatrix} \tilde{A}_1 \\ \tilde{A}_2 \end{pmatrix})$ of the matrix $\begin{pmatrix} \tilde{A}_1 \\ \tilde{A}_2 \end{pmatrix}$ is also fixed. The key theorem in the theory of n -fold integer programming states that in this situation the size of $\mathcal{G}(\tilde{A}^{(n)})$ is polynomial in n and it can be computed in time polynomial in n (see Theorem 4.4, Chapter 4, in [18]).

If f is arbitrary convex, one cannot avoid computing the entire Graver basis, which is polynomial in n but of large degree, and use the algorithm in [13]. However, due to (f_3) being linear, it is possible to use the drastically stronger result of [15]. We briefly explain how this works. To show that this scheme is efficient and polynomial, it is necessary to check that both the number of augmentations and the computational effort for each augmentation is polynomial in the input.

For the number of augmentations, it can be shown that if one performs so-called *Graver-best augmentation steps*, i.e. one always uses a best combination of g and α to improve on the current feasible solution, one then obtains an upper bound $O(nL)$ on the number of necessary steps, where L is the bit size of the input, that is, a bound which is linear in n and in L (see the proof of Lemma 3.10, Chapter 4, in [18]).

To find the Graver-best augmentation step at each iteration, the first approach, taken in [14], was to compute the entire Graver basis, which, as said, has size polynomial in n but of large degree. However, in [15] it was shown that using a sophisticated dynamic program in each iteration, it is possible to find a Graver-best augmentation step at that iteration *without* explicitly constructing the Graver basis, in time $O(n^2)$. This leads to a very fast and practically implementable total running time $O(n^3L)$ for augmenting the initial feasible solution to an optimal one.

Algorithm 1 Polynomial-Time Land Consolidation

Input

- p farmers, and for each farmer i
 - accepted lower and upper bounds $B_i^\pm \in \mathbb{R}^s$ with respect to s features of the lots
 - the geographical location $v_i \in \mathbb{R}^2$ of a farmstead
- n lots, and for each lot j
 - a matrix $W^j = (W_1^j, \dots, W_p^j) \in \mathbb{R}^{s \times p}$ of contributions $W_i^j \in \mathbb{R}^s$ with respect to all s features if lot j is assigned to farmer i . ω_j denotes the size of lot j . The W_i^j come from a finite, fixed domain Ω .
 - its geographical location $z_j \in \mathbb{R}^2$
- Original land distribution represented by $y \in \{0, 1\}^{p \times n}$ with $y_{ij} = 1$ if lot j belongs to farmer i and $y_{ij} = 0$ otherwise. κ_i denotes the original total size of all lots of farmer i , and is derived from y and the ω_j .

We denote by L the input size which is the total number of bits in the binary encoding of the data $(B_i^\pm, v_i, \omega_j, z_j, y)$. The matrices W^j are fixed and hence do not have to be considered for L .

Output

- Partition of the lots into p clusters π_1, \dots, π_p that is optimal with respect to (f_3) , i.e. an optimal solution for

$$\min \sum_{i=1}^p \frac{1}{\kappa_i} \sum_{j \in \pi_i} \omega_j \|v_i - z_j\|^2$$

that adheres to the accepted lower and upper bounds

Algorithm

1. Use slack vectors $S_i^\pm \in \mathbb{Z}^s$, $f : \mathbb{R}^p \rightarrow \mathbb{R}$ with $f(y) = (1, \dots, 1)^T y$, and matrix $C = (C_1^j, \dots, C_p^j) \in \mathbb{R}^{1 \times p}$ of utility $C_i^j = -\frac{\omega_j}{\kappa_i} \|v_i - x_j\|^2$ of lot j to describe the problem in the form (P_3) (see Section 2)
2. Transform the formulation (P_3) to the form (P^*) (see the proof of Theorem 3)

$$\max\{f(\tilde{C}\tilde{x}) : \tilde{A}^{(n)}\tilde{x} = \tilde{b}', 0 \leq \tilde{x} \leq \tilde{u}, \tilde{x} \in \mathbb{Z}^{\tilde{N}}\}$$

3. Construct a feasible solution \tilde{x}_* for (P^*) from y (see the proofs of Theorems 1 and 3)
4. Augment the initial feasible solution \tilde{x}_* to an optimal one \tilde{x}^* by iteratively using Graver-best augmentation steps using the algorithm of [15], in total time $O(n^3 L)$.
5. Construct a decision variable vector $y^* \in \{0, 1\}^{p \times n}$ that corresponds to the optimal \tilde{x}^* and return y^* (see the proofs of Theorems 1 and 3)

Acknowledgements

The first author gratefully acknowledges support from the Alexander-von-Humboldt Foundation. The second author was partially supported by the Dresner Chair at the Technion.

References

- [1] A. Alpers, A. Brieden, P. Gritzmann, and H. F. Poulsen. Generalized balanced power diagrams for 3D representations of polycrystals. *Philosophical Magazine*, 95(9):1016–1028, 2015.
- [2] E. R. Barnes, A. J. Hoffman, and U. G. Rothblum. Optimal partitions having disjoint convex and conic hulls. *Mathematical Programming*, 54(1):69–86, 1992.
- [3] S. Basu, I. Davidson, and K. L. Wagstaff. *Clustering with Constraints: Advances in Algorithms, Theory and Applications*. Chapman & Hall, 2009.
- [4] K. Borgwardt, S. Borgwardt, A. Feragen, and N. Shervashidze. Balanced kernel k-means for comparing large graphs with landmarks. Technical report, 2012.
- [5] S. Borgwardt. *A Combinatorial Optimization Approach to Constrained Clustering*. 2010. PhD Thesis.
- [6] S. Borgwardt, A. Brieden, and P. Gritzmann. Constrained minimum- k -star clustering and its application to the consolidation of farmland. *Operational Research*, 11(1):1–17, 2011.
- [7] S. Borgwardt, A. Brieden, and P. Gritzmann. Geometric clustering for the consolidation of farmland and woodland. *The Mathematical Intelligencer*, 36(2):37–44, 2014.
- [8] A. Brieden. *On the approximability of (discrete) convex maximization and its contribution to the consolidation of farmland*. 2003. Habilitationsschrift.
- [9] A. Brieden and P. Gritzmann. A quadratic optimization model for the consolidation of farmland by means of lend-lease agreements. In *Operations Research Proceedings 2003: Selected Papers of the International Conference on Operations Research*, pages 324–331, 2004.
- [10] A. Brieden and P. Gritzmann. On clustering bodies: Geometry and polyhedral approximation. *Discrete Computational Geometry*, 44(3):508–534, 2010.
- [11] A. Brieden and P. Gritzmann. On optimal weighted balanced clusterings: Gravity bodies and power diagrams. *SIAM Journal on Discrete Mathematics*, 26:415–434, 2012.
- [12] J. A. De Loera, R. Hemmecke, and M. Köppe. *Algebraic and geometric ideas in the theory of discrete optimization*, volume 14 of *MOS-SIAM Series on Optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2013.
- [13] J. A. De Loera, R. Hemmecke, S. Onn, U. G. Rothblum, and R. Weismantel. Convex integer maximization via Graver bases. *Journal of Pure and Applied Algebra*, 213:1569–1577, 2009.

- [14] J. A. De Loera, R. Hemmecke, S. Onn, and R. Weismantel. N -fold integer programming. *Discrete Optimization*, 5(2):231–241, 2008.
- [15] R. Hemmecke, S. Onn, and L. Romanchuk. N -fold integer programming in cubic time. *Mathematical Programming*, 137:325–341, 2013.
- [16] F. Höppner and F. Klawonn. Clustering with size constraints. *Computational Intelligence Paradigms, Innovative Applications*, 2008.
- [17] F. K. Hwang, S. Onn, and U. G. Rothblum. A polynomial time algorithm for shaped partition problems. *SIAM Journal on Optimization*, 10(1):70–81, 1999.
- [18] S. Onn. *Nonlinear Discrete Optimization*. Zurich Lectures in Advanced Mathematics. European Mathematical Society, available online at <http://ie.technion.ac.il/~onn/Book/ND0.pdf>, 2010.
- [19] S. Onn. Theory and applications of N -fold integer programming. In *The IMA Volumes in Mathematics and its Applications, Mixed Integer Nonlinear Programming*, pages 559–593. Springer, 2012.
- [20] S. Onn and U. G. Rothblum. Convex combinatorial optimization. *Discrete & Computational Geometry*, 32:549–566, 2004.